



INVESTOR IN PEOPLE

CERTIFIED COPY OF PRIORITY DOCUMENT

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1985 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, L.C. or PLC.


Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated

31 October 2000

Jc987 U.S. PTO
09/761600



THIS PAGE BLANK (USPTO)

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

THE PATENT OFFICE
J
19 JAN 2000
NEWPORT

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference C1449

2. Patent application number
(The Patent Office will fill in this part)

0001099.1

19 JAN 2000

3. Full name, address and postcode of the or of each applicant (underline all surnames)

INTERNATIONAL COMPUTERS LIMITED
26 Finsbury Square, London EC2A 1SL

Patents ADP number (if you know it)

282 58 008

If the applicant is a corporate body, give the country/state of its incorporation

ENGLAND

4. Title of the invention

Computer-Implemented Simulation Method and Apparatus

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent
(including the postcode)

D C Guyatt
Intellectual Property Department
International Computers Limited
Stevenage
Herts
SG1 2DY

Patents ADP number (if you know it)

1094 937 005

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

YES

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 20

Claim(s) 4

Abstract 1

Drawing(s) 7

10. If you are also filing any of the following, state how many against each item.

Priority documents -

Translations of priority documents -

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 2

Request for preliminary examination and search (Patents Form 9/77) 1

Request for substantive examination (Patents Form 10/77) -

Any other documents Fee Sheet
(please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature

Date 18/01/2000

12. Name and daytime telephone number of person to contact in the United Kingdom

D. L. GUYATT
01438 786 235

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

19 JAN 2000

NEWPORT

Statement of inventorship and of right to grant of a patent

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

C1449

2. Patent application number

(if you know it)

0001099.1

19 JAN 2000

3. Full name of the or of each applicant

INTERNATIONAL COMPUTERS LIMITED

4. Title of the invention

Computer-Implemented Simulation Method and Apparatus

5. State how the applicant(s) derived the right from the inventor(s) to be granted a patent

By reason of employment

6. How many, if any, additional Patents Forms

7/77 are attached to this form?

(see note (c))

7.

I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.

Signature

D C Guyatt

Date

18/01/2000

8. Name and daytime telephone number of person to contact in the United Kingdom

D C Guyatt
01438 786235

Notes

a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.

b) Write your answers in capital letters using black ink or you may type them.

c) If there are more than three inventors, please write the names and addresses of the other inventors on the back of another Patents Form 7/77 and attach it to this form.

d) When an application does not declare any priority, or declares priority from an earlier UK application, you must provide enough copies of this form so that the Patent Office can send one to each inventor who is not an applicant.

e) Once you have filled in the form you must remember to sign and date it.

THIS PAGE BLANK (USPTO)

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Zakwan Shaar
7 Boddens Hill Road
Heaton Mersey
Stockport
SK4 2DG

Patents ADP number (if you know it): 2073508001

Patents ADP number (if you know it):

Reminder

Have you signed the form?

Patents ADP number (if you know it):

THIS PAGE BLANK (USPTO)

COMPUTER-IMPLEMENTED SIMULATION METHOD AND APPARATUS

Background to the Invention

This invention relates to computer-implemented simulation techniques. The invention is particularly, although not exclusively, concerned with simulating the design of computer hardware; it could also be used in simulating other systems, such as complex software systems.

Design specification languages are known, which allow the design of a target system to be specified in terms of a number of functional units that communicate with each other by means of signals. One such language is VHDL [ANSI/IEEE Standard 1076]. The model can then be executed on a VHDL simulator, to simulate the behaviour of the target system.

VHDL+ is a superset of VHDL that extends the capability of the language up to the system design level. In particular, VHDL+ extends the VHDL language by introducing a new primary unit, the interface, which models an interface between two or more design units. Interfaces contain transactions, messages and signals. For more details of VHDL+, reference is made to <http://www.icl.com/da/vhdl+.htm>.

One object of the invention is to provide an improved simulation method and apparatus, which can facilitate execution of VHDL+ models. It should be noted, however, that the invention is not restricted to VHDL+, and is also potentially useful in relation to other simulation languages such as Verilog and SystemC.

Summary of the Invention

According to the invention a computer-implemented simulation method comprises the steps:

- modelling a target system as a set of processes that communicate with each other by way of messages;
- associating the messages with sender and receiver processes;

- providing at least one scheduler queue, holding a series of items, each item having a type value which indicates the item type;
- scheduling the processes and messages by placing process-type items and message-type items on the scheduler queue;
- processing each process-type item on the scheduler queue by calling the process to which the item relates; and
- processing each message-type item on the scheduler queue by calling both the sender and receiver processes of the message to which the item relates.

Brief description of the drawings

Figure 1 is a block schematic diagram of a simulation system, including a simulation model.

Figure 2 is a block schematic diagram showing possible ways in which the simulation model can be generated.

Figures 3-9 are flow charts showing various processes in the kernel of the simulation model.

Figures 10 and 11 are schematic block diagrams of parts of a simulation model, illustrating composition and decomposition of messages.

Description of an embodiment of the invention

One simulation system in accordance with the invention will now be described by way of example with reference to the accompanying drawings.

Messages and processes

The simulation system to be described is based on a message-process model.

Processes represent specific functional units of the target system (i.e. the particular system that is being simulated). Processes are permanently active during the running of the simulation.

Messages are the basic units of communication between processes. A message is regarded as a resource that is shared between one or more senders and one or more receivers, and is modelled as a data structure (class) with the following attributes.

<i>to</i>	pointer to interface receiving end
<i>from</i>	pointer to interface sending end
<i>state</i>	message state
<i>takes</i>	time (in nanoseconds) taken to send the message
<i>queued</i>	message is queued on the event or delta queue
<i>start_t</i>	start time
<i>compose</i>	pointer to composition activity
<i>decompose</i>	pointer to decomposition activity
<i>composing</i>	composition is active
<i>decomposing</i>	decomposition is active
<i>snd_cnt</i>	send count
<i>rcv_cnt</i>	receive count
<i>data[i]</i>	message data of <i>i</i> th sender
<i>sender[i]</i>	pointer to <i>i</i> th sender process
<i>rcv_req_t[i]</i>	receive request time of <i>i</i> th receiver process
<i>receiver[i]</i>	pointer to <i>i</i> th receiver process.

The arrays *data[i]* and *sender[i]* form a sender queue, in which processes waiting to send the message can be queued, with *snd_cnt* pointing to the first free location in this queue. Similarly, the arrays *rcv_req_t[i]* and *receiver[i]* form a receiver queue, in which processes waiting to receive the message can be queued, with *rcv_cnt* pointing to the first free location in this queue.

Communication is deemed to commence when a message has at least one sender and at least one receiver, whereupon the message becomes busy for the duration specified by *takes*. Any sender or receiver arriving during this busy period will be queued and will not form part of the current message-process activation sequence.

Activities

An activity represents a special type of process that terminates. An activity is regarded as a shared resource between contending caller processes. Each activity is modelled as a data structure, with an associated activity process. The data structure has the attributes:

<i>state</i>	activity state
<i>queued</i>	the associated activity process is scheduled on the event or delta queue
<i>start_t</i>	activity start time
<i>call_cnt</i>	caller count
<i>proc</i>	pointer to the associated activity process
<i>caller[i]</i>	pointer to <i>i</i> th caller process.

The array *caller[i]* forms a caller queue, for queuing caller processes waiting to be served by the activity, with *call_cnt* pointing to the first free location of this queue. An activity must terminate before it can serve the next process in its caller queue.

Composition and decomposition

Different processes in the simulation model may be specified at different levels of design. For example, one process may expect to receive and send a byte-level message, while another process may expect to receive and send a bit-level message. Composition is the process of constructing a higher level message from a series of lower level messages. Decomposition is the reverse

process of expressing a higher level message as a series of lower level messages.

As described above, each message data structure contains pointers (possibly NULL) to composition and decomposition activities for that message. The decomposition activity (if any) is triggered when the message is sent, and the composition activity (if any) is triggered when the message is received.

Signals

Signals are the basic units of communication between VHDL design unit entities (acting as ports) and VHDL processes of an entity architecture. They are data types, typically used for passing hardware-related low-level information such as bits, bit vectors, registers, buses, etc.

An important feature of the present system is that messages and signals are both handled by the same kernel, using the same scheduling mechanism. This avoids the need for two separate kernels, which would require considerable synchronisation overheads.

Although VHDL signals are mentioned by way of example, it will be appreciated that the signal concept is generic and can be implemented in any language.

Events and deltas

The simulation system uses two forms of simulated time: event time and delta time. Event time is measured in units such as nanoseconds. Delta time is simply an integer that is incremented each time the delta queue is processed.

The system uses two scheduler queues to schedule changes to the state of the model: an event queue for scheduling changes at

specified event times, and a delta queue for scheduling changes that are assumed to occur intantaneously.

An event is a data structure, containing the following values:

<i>type</i>	the event type
<i>id</i>	a pointer to a message, process, activity or signal that is to be scheduled.
<i>next_event</i>	a pointer to the next event in the queue
<i>time</i>	the event time at which the event is to be activated.

In the present embodiment, the event type may be one of the following: MESSAGE, PROCESS, ACTIVITY or SIGNAL. Events are handled in different ways according to their type:

- PROCESS or ACTIVITY events are handled by calling the process or activity pointed to by *id*.
- MESSAGE events are handled by calling both the sender and the receiver processes of the message pointed to by *id*.
- SIGNAL events are handled by first updating the signal state and then calling all the processes sensitive to the signal.

A delta is similar to an event, except that it does not contain any *time* value, and has *next_delta* instead of *next_event*. As with events, deltas are handled in different ways according to their type.

Once processed, an event (or delta) is returned to a free event list (or free delta list) for reuse, removing the overhead of constant creation and deletion. A new event or delta is created only when the corresponding free list is empty.

The event queue is based on a time loop having a number of time slots, as described in Patent No. GB2219419. The events are mapped on to the time slots in such a way that events with event times corresponding to different cycles around the loop may be simultaneously present in the same time slot. Events that are

mapped into the same time slot are chained together as a linked list, in chronological order. Reference is made to the above mentioned patent for more details. The delta queue is a simple linear list.

Overview of the system

Referring to Figure 1, the simulation system comprises a computer 10, which runs a simulation model 11. The simulation model 11 comprises a Kernel 12, model processes 13, data structures 14, and a command language interpreter 15.

The Kernel 12 is an event-driven process-oriented simulation engine. It comprises scheduling queues (an event queue and a delta queue), a set of scheduling routines, and a set of service routines. These routines are described in detail below.

The processes 13 represent the components of the particular system that is being modelled, and are specific to that system. They are modelled as concurrent state machines, which are re-entered when message passing takes place. In operation, these processes make calls to the Kernel routines, for example to pass messages between processes, and to start and finish activities.

The data structures 14, correspond to the messages, activities, signals and process variables of the model.

The command language interpreter 15 processes user input and output, and manages a number of graphical user interface (GUI) tools 16. The interpreter may, for example, comprise a Tcl/Tk language interpreter. The use of Tcl renders the simulator an open system, capable of integrating and being integrated with other applications.

Figure 2 shows the way in which the simulation model 11 can be generated from a VHDL+ source file 20, defining the simulation

model in terms of a number of processes, activities messages and signals. First, the VHDL+ source file is used to generate a set of C or C++ (.c) source files, which directly model the semantics of the VHDL+. The generation of the C/C++ files may be performed by hand coding, or by means of an automatic code generation program 21. The C/C++ files are then compiled, by means of a C/C++ compiler 22, to produce a set of shared dynamic (.so) library files 23. The shared library files are then loaded, along with the simulator kernel 24 and library files 25 for the command language interpreter, to form the simulation model.

Alternatively, the compiler 22 may be arranged to generate a set of compiled objects (.o files) 26 which can then be linked with the kernel 24 by means of a linker 27, to form a single executable program 28. This approach is useful for smaller models and for prototyping.

The scheduling and service routines in the Kernel 12 will now be described.

push_event (type, id, time)

Referring to Figure 3, this routine adds a new event to the event queue to schedule a message, process, activity or signal. The call parameters are:

type the event type (MESSAGE, PROCESS, ACTIVITY or SIGNAL)
id a pointer to the message, process, activity or signal in question
time the event time at which the event is to be scheduled.

(Step 30) The routine first checks whether there are any free events in the free event list.

(Step 31) If there are no free events, a new event data structure is created and allocated to the new event.

(Step 32) Alternatively, if there is a free event, it is removed from the free list and allocated to the new event.

(Step 33) The *type*, *id*, and *time* values in the event are then set to the corresponding values from the call parameters, while the *next_event* value is set to NULL.

(Step 34) The time slot for the event is then calculated as:

time_slot = (*time* & LOOPBOUND),

where LOOPBOUND is the size of the time loop, and "&" represents a bitwise AND.

(Step 35) The routine then compares the event time for the new event with the earliest event time in the list of events for this time slot.

(Step 36) If they are equal, the new event is linked into the list at a position immediately following last of the events with this event time.

(Step 37) If the new event time is earlier than the earliest event time, the new event is added to the beginning of the list.

(Step 38) If the new event time is later than the earliest event time, the list is searched to find the first event with a later event time, and the new event is inserted in that position. If there is no event with a later event time, the new event is inserted at the end of the list.

push_delta (type, id)

This routine adds a new delta to the delta queue. The routine first checks whether there are any free deltas in the free delta list. If there is no free delta, a new data structure is

created and allocated to the new delta; otherwise, an existing free delta is reused. The new delta is then added to the end of the delta queue. The *type* and *id* values in the delta are set to the corresponding values from the call parameters, and the *next_delta* value is set to NULL.

pop_events()

Referring to Figure 4, this routine removes events from the event queue, in the correct sequence and at the correct event times, and process them.

(Step 40) The routine first calls *pop_deltas*, to process the delta queue.

(Step 41) It then enters a loop (steps 41-49), which is repeated until *event_time* (the current event time) is greater than *simtime* (a predetermined end time for the simulation). When this happens, the process returns a TRUE value, to indicate that the simulation is complete.

(Step 42) The time slot corresponding to the current event time is calculated, as described in step 34.

(Step 43) The current event time is then compared with the event time of the earliest event in the list for this time slot. If they are equal, the process goes to step 44; otherwise, it goes to step 48.

(Step 44) The routine processes all the events in this time slot, with event times equal to *event_time*. The processing of an event depends on its type: MESSAGE events are processed by calling the senders and receivers of the message, while PROCESS and ACTIVITY events are processed by calling the relevant process or activity. SIGNAL events are processed by updating the

signal state. The processing of an event is described in more detail below with reference to Figure 5.

(Step 45) If any SIGNAL events were processed in Step 44, a second pass is now performed, to call all processes sensitive to those signals.

(Step 46) The processed events are returned to the free event list.

(Step 47) Pop_deltas is called, to process any deltas that have arrived since last time this was called.

(Step 48) *Event_time* is incremented. The routine then loops back to step 41.

(Step 49) The routine checks whether it has gone once round the timeloop without finding any events to process. This can arise when there are long delays between events.

(Step 410) If so, *event_time* is updated by setting it to the earliest event time encountered in the loop.

Figure 5 shows the processing of an event.

(Step 50) A branch is made according to event type.

(Step 51) If the event type is MESSAGE, the first process on the sender queue for the message is called, and the send count *snd_cnt* is decremented by one.

(Step 52) All processes on the receiver queue that arrived in time are then called. A receiver is considered to have arrived in time if its receive request time *rcv_req_t[i]* is less than or equal to the message start time *start_t*. As each receiver

process is called, the message state is set to RECEIVED, and its receive count *rcv_cnt* is decremented by one.

(Step 53) Any remaining senders and receivers are then moved up their respective sender and receiver queues.

(Step 54) A check is then made to determine whether there is at least one remaining sender and receiver in the sender and receiver queues (i.e. *rcv_cnt* and *snd_cnt* are both greater than zero).

(Step 55) If there is at least one remaining sender and receiver, the message is rescheduled to allow these to be processed:

- set *start_t* = *event_time*
- set the message state to ACTIVE
- call *push_event*, to reschedule the message on the event queue, with *time* = *takes* + *event_time*.

(Step 56) If on the other hand either the sender queue or the receiver queue is empty, *queued* is set to FALSE.

(Step 57) If the event type is PROCESS or ACTIVITY, the process or activity is called.

(Step 58) If the event type is SIGNAL, the signal state is updated.

pop_deltas ()

This routine first checks whether there are any deltas on the delta queue. If there are, it increments the value of *delta_time*, and then processes all the deltas on the queue. The processing of a delta is similar to the processing of an event, as described with reference to Figure 5. In particular, MESSAGE deltas are processed by calling the senders and receivers of the

message, while PROCESS and ACTIVITY deltas are processed by calling the relevant processes or activities. SIGNAL deltas are processed by updating the state and then, in a second pass, calling sensitive processes.

send_m (m, proc, data)

Referring to Figure 6, this routine adds a sender to a message. The call parameters are:

<i>m</i>	a pointer to the message data structure
<i>proc</i>	a pointer to the sender process
<i>data</i>	the message data.

(Step 60) The routine first increments the send count (*snd_cnt*) by one.

(Step 61) It then tests whether the sender queue is full. If so, it prints an error message, and exits. Otherwise, it continues with the following steps.

(Step 62) The routine adds the sender process and the message data to the sender queue.

(Step 63) It then tests whether the message can now be scheduled. This is the case if *decompose* is NULL, *queued* is FALSE, and *rcv_cnt* is greater than zero; that is, there is no decomposition for the message, the message is not currently queued, and there is at least one receiver in the receive queue. If the message cannot be scheduled, the routine goes to Step 65; the message will be scheduled later, e.g. by the *receive_m* routine.

(Step 64) The message is scheduled, as follows:

- *queued* = TRUE
- *state* = ACTIVE
- *start_t* = *event_time*

- if *takes* is zero, call *push_delta* to schedule the message on the delta queue.
- if *takes* is greater than zero, call *push_event* to schedule the message on the event queue, with *time = takes + event_time*.

(Step 65) The message state is set to *SEND_REQ*.

(Step 66) If a decomposition activity is specified for the message, the decomposition activity is initiated by way of the *call_act* routine, and *decomposing* is set to *TRUE*.

receive_m (m, proc)

Referring to Figure 7, this routine adds a receiver to a message. The call parameters are:

m a pointer to the message data structure
proc a pointer to the receiver process.

(Step 70) The routine first increments the receive count (*rcv_cnt*) by one.

(Step 71) It then tests whether the receiver queue is full. If so, it prints an error message, and exits. Otherwise, it continues with the following steps.

(Step 72) The routine adds the receiving process and the current event time to the receiver queue.

(Step 73) It then tests whether the message can now be scheduled. This is the case if *compose* is *NULL*, *queued* is *FALSE*, and *snd_cnt* is greater than zero; that is, there is no composition for the message, the message is not currently queued, and there is at least one sender in the sender queue. If the message cannot be scheduled, the routine goes to Step 75; the message will be scheduled later, e.g. by the *send_m* routine.

(Step 74) The message is scheduled, as in Step 64.

(Step 75) The message state is set to `RECEIVE_REQ`.

(Step 76) If any composition activity is specified for the message, the composition activity is initiated by means of the `call_act` routine, and *composing* is set to `TRUE`.

`call_act (a, proc)`

Referring to Figure 8, this routine initiates an activity. The call parameters are:

<code>a</code>	a pointer to the activity data structure
<code>proc</code>	a pointer to the caller process

(Step 80) The routine first increments the call count `call_cnt` by one.

(Step 81) It then tests whether the caller queue is full. If so, it prints an error message, and jumps to step 85. Otherwise, it continues with the following steps.

(Step 82) The routine sets the activity start time (`start_t`) equal to the current event time, and adds the caller process to the caller queue.

(Step 83) It then tests the activity's *queued* attribute to determine whether the activity is already queued.

(Step 84) If the activity is not already scheduled, it is now scheduled, as follows:

- `queued = TRUE`
- call `push_delta` to schedule the activity process on the delta queue.

(Step 85) Finally, the activity mode is set to `STARTED`.

finish_act (a)

Referring to Figure 9, this routine finishes an activity.

(Steps 90-92) The routine calls the first caller process in the caller queue. It then decrements the activity call count (*call_cnt*) by one and sets the activity mode to FINISHED. It then checks whether the call count is greater than zero, i.e. whether there are any remaining callers in the caller queue.

(Step 93) If there are no remaining callers, *queued* is set to FALSE and the routine exits.

(Steps 94-96) If there are any remaining callers, they are moved one place up the caller queue. The routine then calls *push_delta*, to reschedule the activity process on the delta queue. Finally, the routine sets the activity mode to STARTED.

finish_c_act (a, m)

This routine is used to finish a composition or decomposition activity associated with a message. The routine is similar to the *finish_act* routine described above, except that step 90 is replaced by the following step.

If the *queued* attribute of the message is FALSE, the routine sets it to TRUE, and then calls *push_delta*, to schedule the message on the delta queue.

Model Processes

As mentioned above, the simulation model 11 (Figure 1) includes a number of processes 13 specific to the target system. The processes 13 are modelled as concurrent state machines, which are re-entered when message passing takes place. The processes

are concurrent, and their order of execution is random. The processes are called by the Kernel without any parameters, keeping stack size low, which improves performance.

The processes 13 communicate with each other only by way of the `send_m` and `receive_m` routines, or by signals, and can launch and terminate activities only through the `start_act` and `finish_act` routines.

A process 13 can also directly call `push_delta` or `push_event`, to schedule its own recall. For example, if a process requires to wait for a specified time, it calls `push_event` to schedule a PROCESS event, with a pointer to itself as the `id` parameter. Alternatively, it can call `push_delta`, to schedule an immediate recall.

Example of Decomposition

Figure 10 shows an example of part of a simulation model, to illustrate message decomposition. The model comprises a sender process 100, which sends a word-level message 101, and a receiver process 102, which receives a byte-level message 103. The word-level message 101 has a decomposition activity 104 associated with it, which in turn has an associated decomposition process 105.

The sender process 100 calls `send_m`, to send the word-level message 101, and then suspends. `Send_m` adds the sender process to the sender queue of message 101, and then calls `call_act` to initiate the decomposition activity 104. `Call_act` then calls the decomposition process 105.

Similarly, the receiver process 102 calls `receive_m`, to receive the byte-level message 103. `Receive_m` adds the receiver process to the receiver queue of the message 103.

The decomposition process 105 in this example is implemented as a state machine with four states (State 0 - State 3); the operation of the process depends on its state when called.

State 0: The decomposition process extracts the next data byte from the word-level message 101, and calls `send_m` to send the byte-level message 103 with this data byte. It then advances to State 1 and suspends. `Send_m` adds the decomposition process as a sender to the sender queue of the byte-level message 103. Assuming that the message 103 now has both a sender and a receiver, this message can now be scheduled.

State 1: The decomposition process calls `push_event`, to reschedule itself after delay of one byte time. The process then advances to State 2, and suspends.

State 2: The decomposition process checks whether it has completed the decomposition (i.e. has processed all 4 bytes from the word). If the decomposition is not complete, it calls `push_delta` to reschedule itself immediately. The process then returns to State 0, and suspends. When recalled, the process will send the next data byte.

If decomposition is complete, the decomposition process calls `receive_m`, to receive the word-level message 101. It then advances to State 3 and suspends. `Receive_m` will add the decomposition process 105 as a receiver to the receiver queue of the word-level message 101. Assuming that the message 101 now has both a sender and a receiver, this message can now be scheduled.

State 3: The decomposition process resets the *decomposing* attribute of the word-level message to FALSE, and calls `finish_c_act`, to finish the decomposition activity. The decomposition process then returns to State 0 and suspends.

Example of Composition

Figure 11 shows an example of part of a simulation model, to illustrate message composition. The model comprises a sender process 110, which sends a byte-level message 111, and a receiver process 112, which receives a word-level message 113. The byte-level message 111 has a composition activity 104 associated with it, which in turn has an associated composition process 115.

The sender process 110 calls `send_m`, to send the byte-level message 111. `Send_m` adds the sender process to the sender queue of message 111.

Similarly, the receiver process 112 calls `receive_m`, to receive the word-level message 113. `Receive_m` adds the receiver process to the receiver queue of the message 113, and then calls `call_act`, to initiate the composition activity 114. `Call_act` then calls the composition process 115.

The composition process 115 in this example is implemented as a state machine with four states (State 0 - State 3).

State 0: The composition process calls `receive_m` to receive the byte-level message 111. The received data byte is accumulated in a variable `rnum`. The process then advances to State 1 and suspends. `Receive_m` adds the composition process as a receiver to the receiver queue of the byte-level message 111. Assuming that the message 111 now has both a sender and a receiver, this message can now be scheduled.

State 1: The composition process calls `push_event`, to reschedule itself after a delay of one byte time. The process then advances to State 2, and suspends.

State 2: The composition process checks whether it has completed the composition (i.e. has received all 4 data bytes of the word). If the composition is not complete, it calls `push_delta` to reschedule itself immediately. It then returns to State 0, and suspends. When recalled, the process will receive the next data byte.

If composition is complete, the composition process calls `send_m`, to send the word-level message 113 with the accumulated data `rnum`. It then advances to State 3 and suspends. `Send_m` adds the composition process 115 as a sender to the sender queue of the word-level message 113. Assuming that the message 113 now has both a sender and a receiver, this message can now be scheduled.

State 3: The composition process resets the *composing* attribute of the word-level message to FALSE, and calls `finish_c_act`, to finish the composition activity. The composition process then returns to State 0 and suspends.

Some possible modifications

It will be appreciated that many modifications may be made to the system described above without departing from the scope of the present invention. For example, the invention may be applied to simulation environments other than VHDL+, and may be implemented using various different programming languages.

Although only four event and delta types (MESSAGE, PROCESS, ACTIVITY and SIGNAL) are described in the above example, other types could also be specified. For example, the VHDL+ TRANSACTION type could be introduced to model multi-directional message passing between two or more ends of the interface.

CLAIMS

1. A computer-implemented simulation method comprising the steps:

- modelling a target system as a set of processes that communicate with each other by way of messages;
- associating the messages with sender and receiver processes;
- providing at least one scheduler queue, holding a series of items, each item having a type value which indicates the item type;
- scheduling the processes and messages by placing process-type items and message-type items on the scheduler queue;
- processing each process-type item on the scheduler queue by calling the process to which the item relates; and
- processing each message-type item on the scheduler queue by calling both the sender and receiver processes of the message to which the item relates.

2. A computer-implemented simulation method according to Claim 1 wherein processes also communicate with each other by way of signals, the method including the further steps of:

- scheduling the signals for processing by placing signal-type items on the scheduler queue; and
- processing each signal-type item on the scheduler queue by updating the signal state and calling all processes that are sensitive to the signal.

3. A method according to Claim 1 or 2, including:

- also modelling the target system as a number of activities;
- providing an activity data structure for each of the activities, each activity data structure including a pointer to an associated activity process and a caller queue for queuing a number of caller processes;
- when a caller process requires to call an activity, adding the caller process to the caller queue of the related activity

data structure, and scheduling the activity for processing by placing an activity-type item on the scheduler queue; and

- processing each activity-type item on the scheduler queue by calling the associated activity process.

4. A method according to Claim 3 wherein the activity is scheduled only if it is not already scheduled for processing.

5. A method according to Claim 3 or 4 wherein, when an activity process finishes its processing:

- the corresponding caller process is removed from the caller queue; and
- if there are still any remaining caller processes in the caller queue, the activity is rescheduled.

6. A method according to any preceding claim wherein there are two scheduler queues: an event queue for scheduling changes that are to take place at specified times, and a delta queue for scheduling changes that are to take place instantaneously.

7. A method according to any preceding claim further including:

- providing a message data structure for each of the messages;
- when a process requires to send a message, adding that process to the relevant message data structure as a sender; and
- when a process requires to receive a message, adding that process to the relevant message data structure as a receiver.

8. A method according to Claim 7 wherein:

- when a process is added as a sender for a message, the message is scheduled for processing only if at least one receiver for the message already exists in the message data structure; and
- when a process is added as a receiver for a message, the message is scheduled for processing only if at least one sender for the message already exists in the message data structure.

9. A method according to Claim 7 or 8 wherein each message data structure includes a sender queue and a receiver queue, for queuing a number of sender and receiver processes.

10. A method according to any one of Claims 7 to 9 wherein, when a message has been processed, the message data structure is checked to determine whether there is at least one remaining sender and receiver for the message and, if so, the message is rescheduled.

11. A method according to any one of Claims 7 to 10 wherein each message data structure holds pointers (possibly NULL) to a composition activity, for composing a higher-level message from a lower-level message, and to a decomposition activity, for decomposing a higher-level message into a lower-level message.

12. A method according to Claim 11 wherein the composition activity performs multiple receives of the lower-level message followed by a single send of the higher-level message.

13. A method according to Claim 11 or 12 wherein the decomposition activity performs multiple sends of the lower-level message followed by a single receive of the higher-level message.

14. A method according to any one of Claims 11 to 13 wherein:

- the decomposition activity for a message is activated when a process is added as a sender for that message; and
- the composition activity for a message is activated when a process is added as a receiver for that message.

15. A computer-implemented simulation method substantially as hereinbefore described with reference to the accompanying drawings.

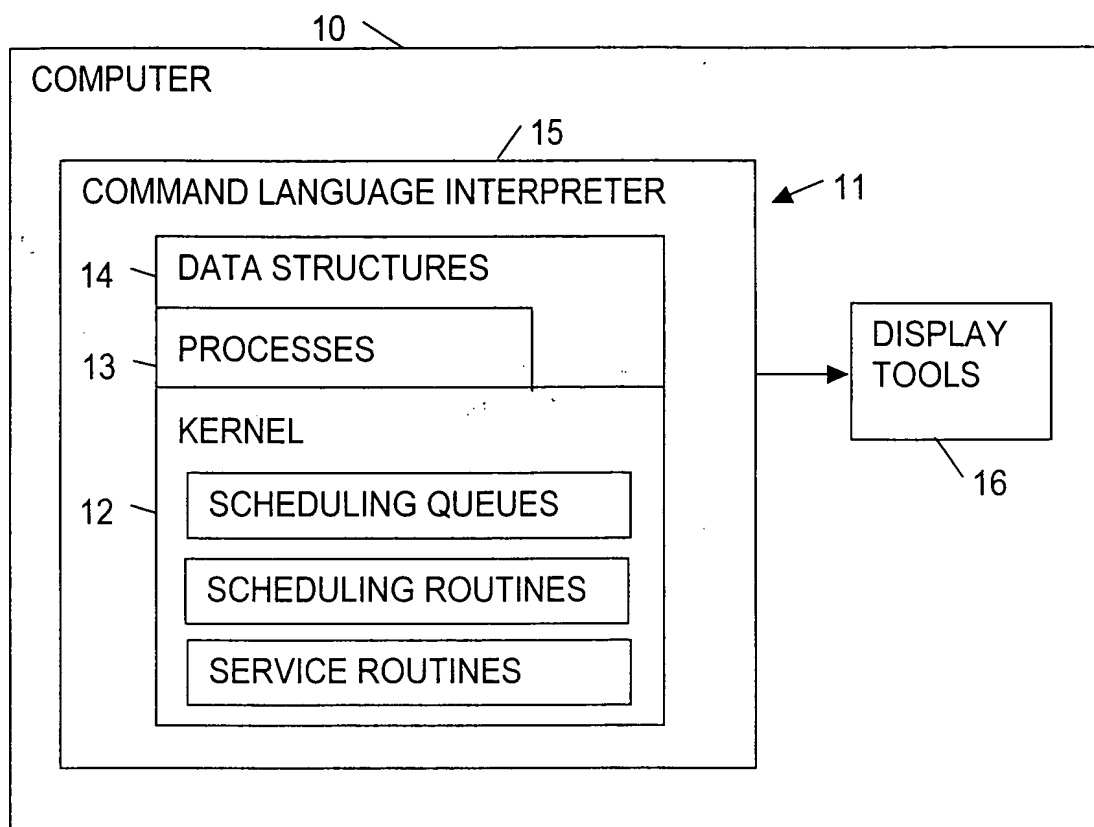
16. A computer system for performing a method in accordance with any preceding claim.

17. A data carrier, carrying a computer-readable program for performing a method in accordance with any preceding claim.

ABSTRACT

A computer-implemented simulation method comprises modelling a target system as a set of processes and activities that communicate with each other by way of messages and signals. Each message is modelled by a data structure containing sender and receiver queues, and pointers to associated composition and decomposition activities. Similarly, each activity is modelled by a data structure containing a caller queue, and a pointer to an activity process. The simulator uses scheduler queues (event and delta queues) to schedule changes to the state of the model. Each item in each scheduler queue has a type value which indicates whether it relates to a process, activity, message or signal, and items are processed in different ways according to their type. A process-type or activity-type item is processed by calling the process or activity to which the item relates. A message-type item is processed by calling both the sender and receiver processes of the message to which the item relates. A signal-type item is processed by updating the signal state and calling all processes that are sensitive to the signal. A message is scheduled only if at least one sender and receiver exist for the message.

THIS PAGE BLANK (USPTO)

**FIG.1**

THIS PAGE BLANK (USPTO)

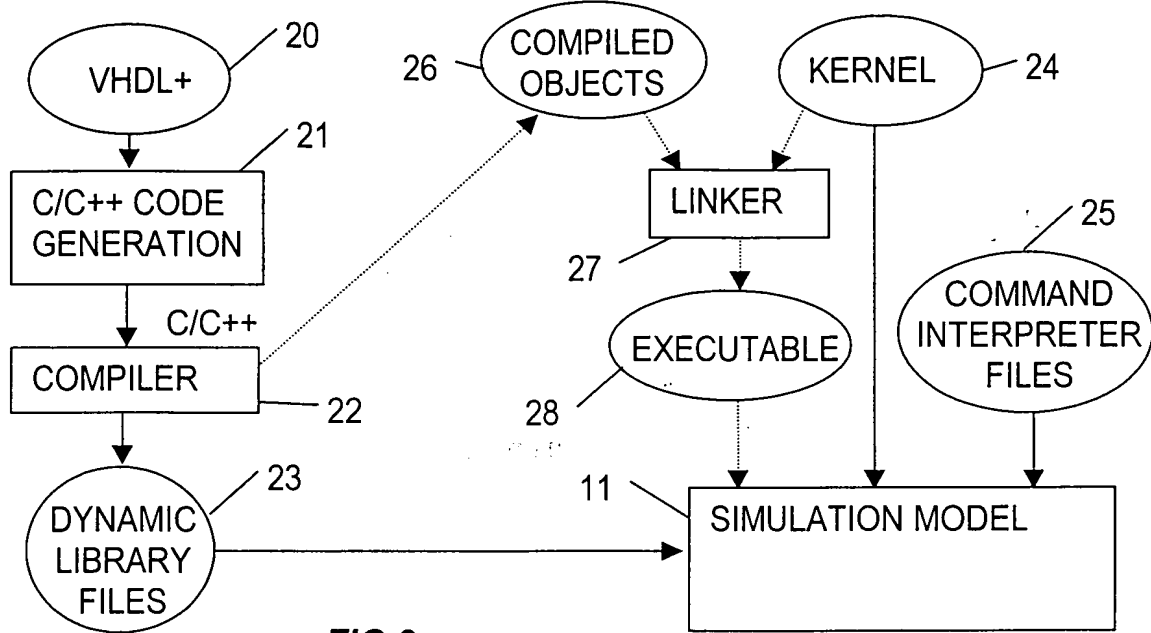
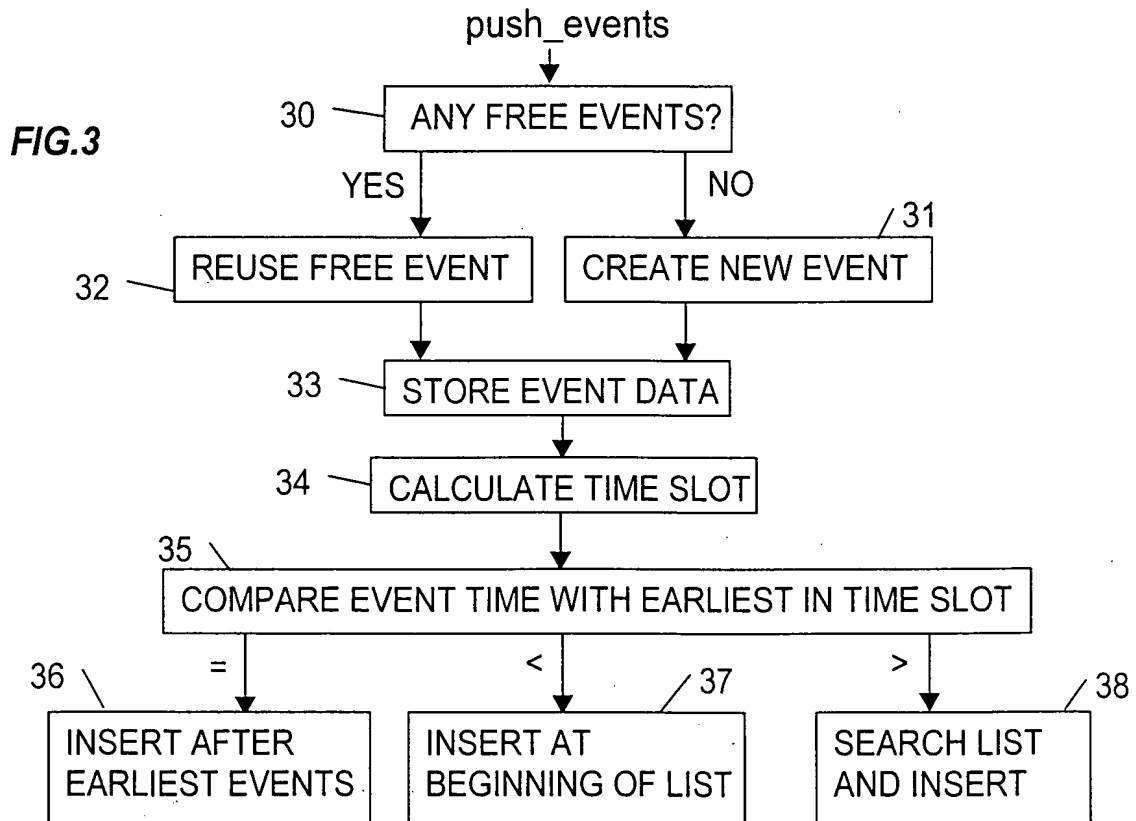
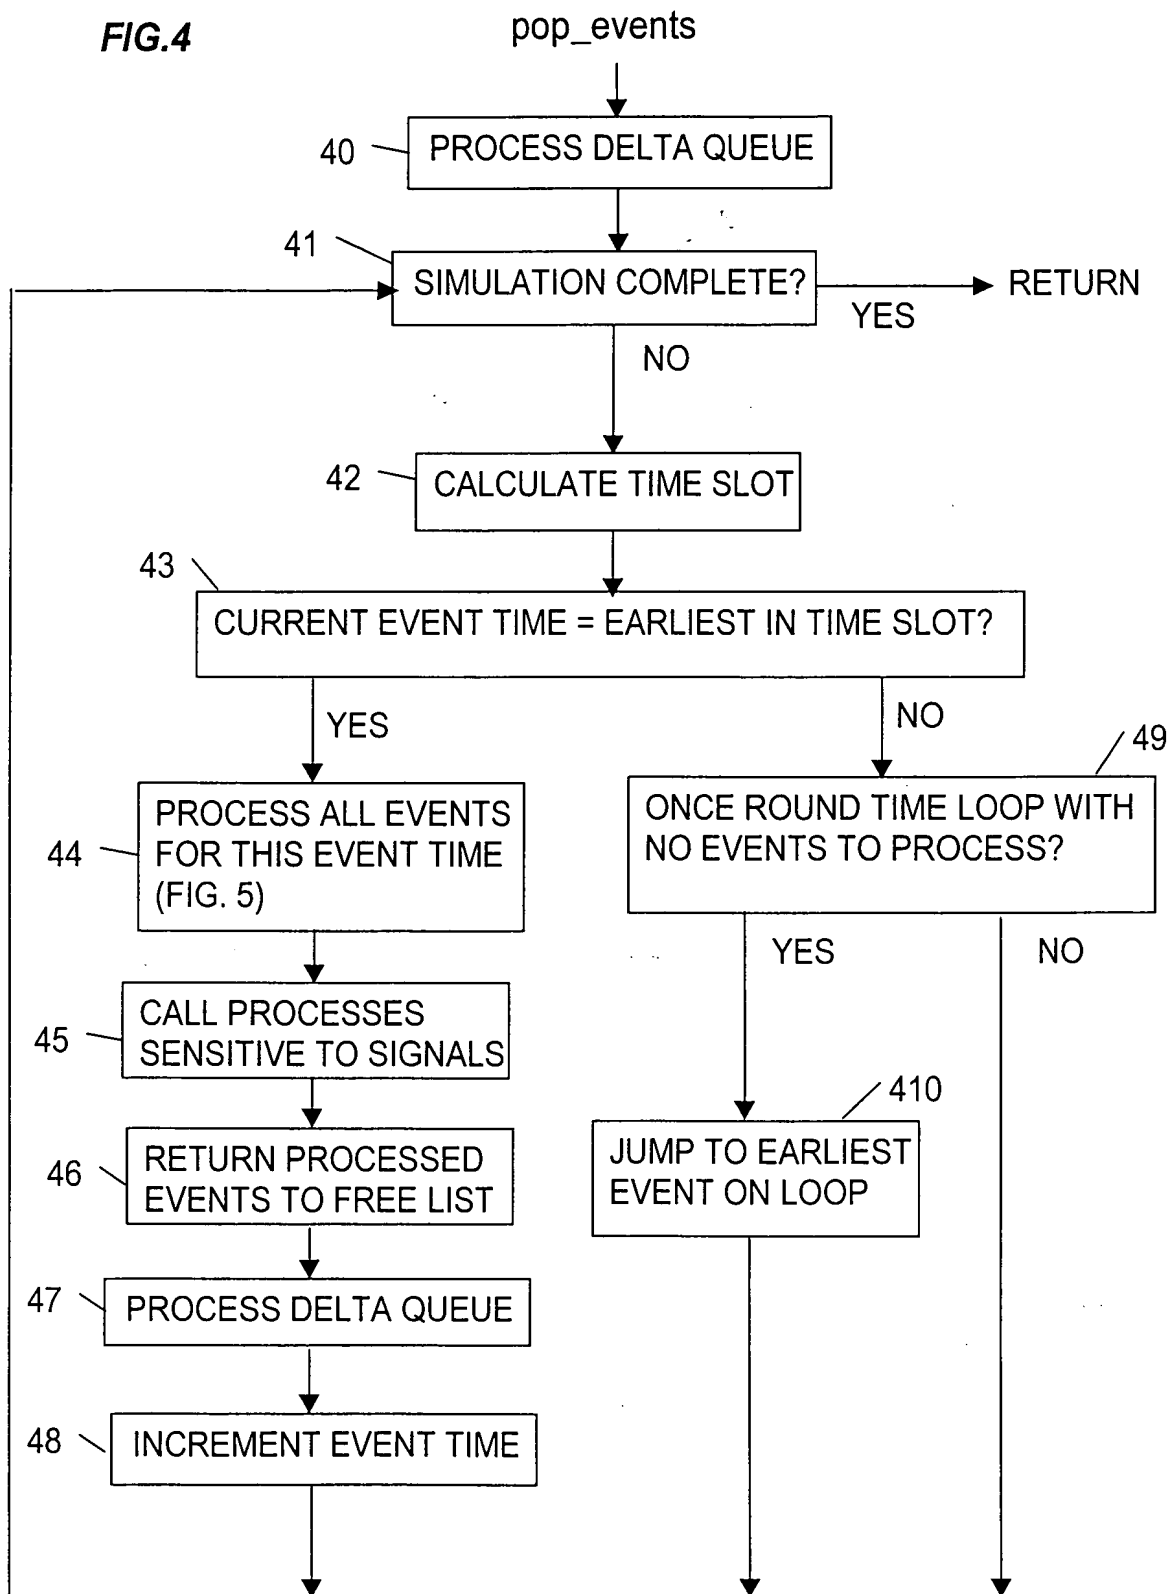


FIG. 2

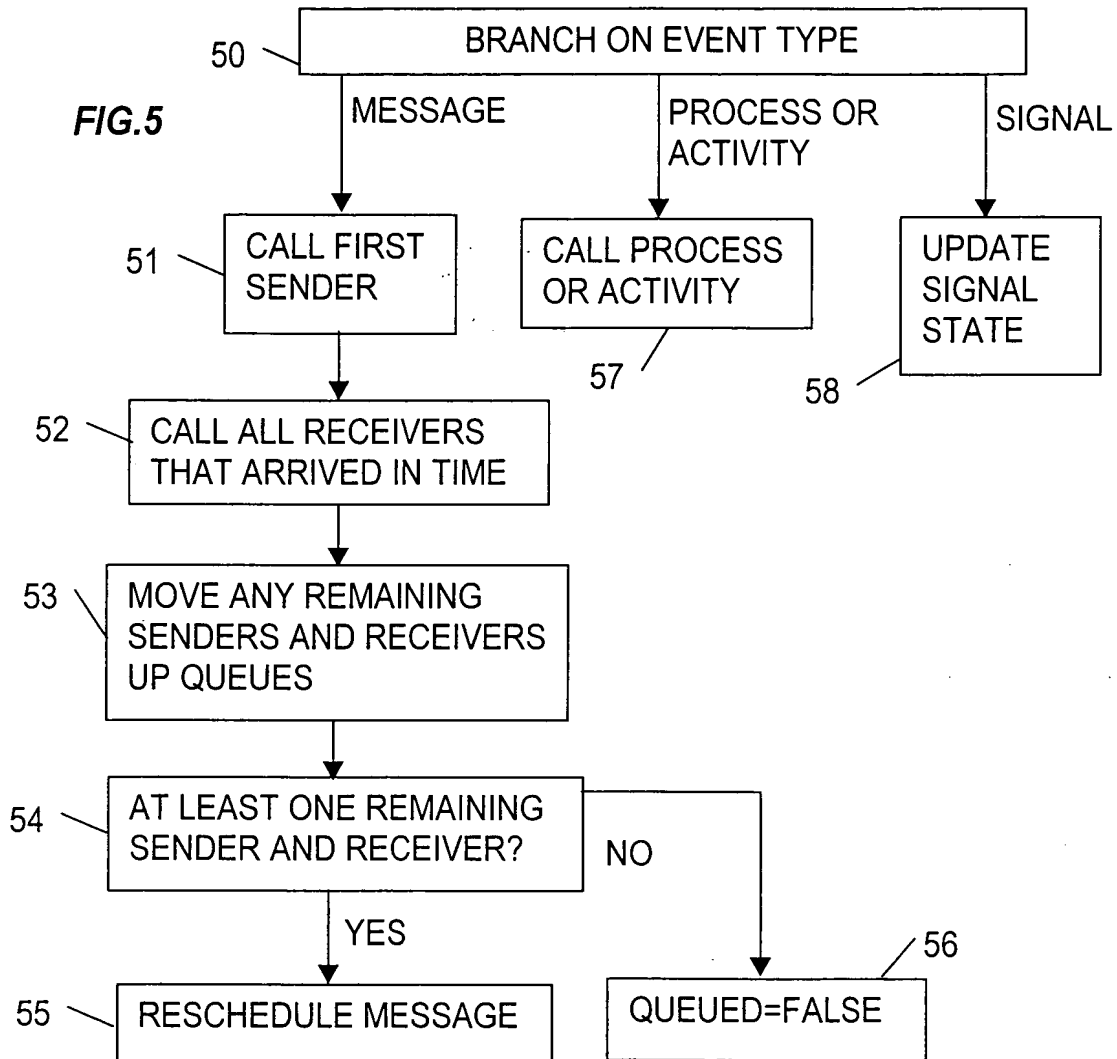


THIS PAGE BLANK (USPTO)

FIG.4



THIS PAGE BLANK (USPTO)



THIS PAGE BLANK (USPTO)

FIG.6

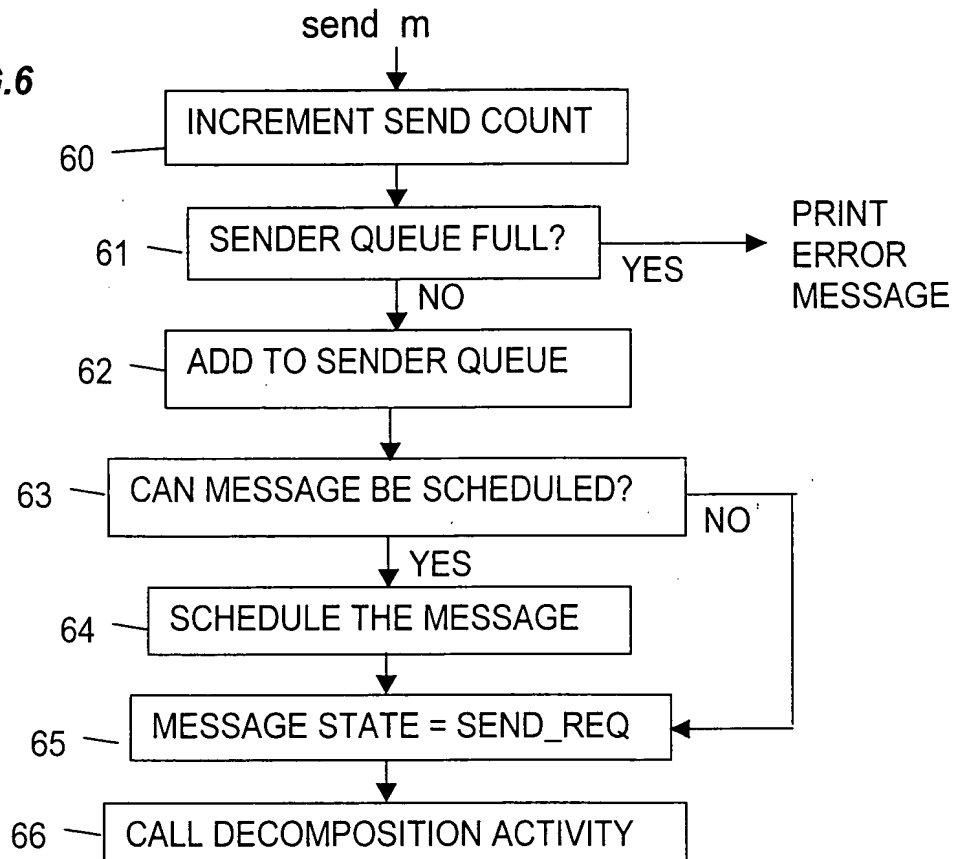
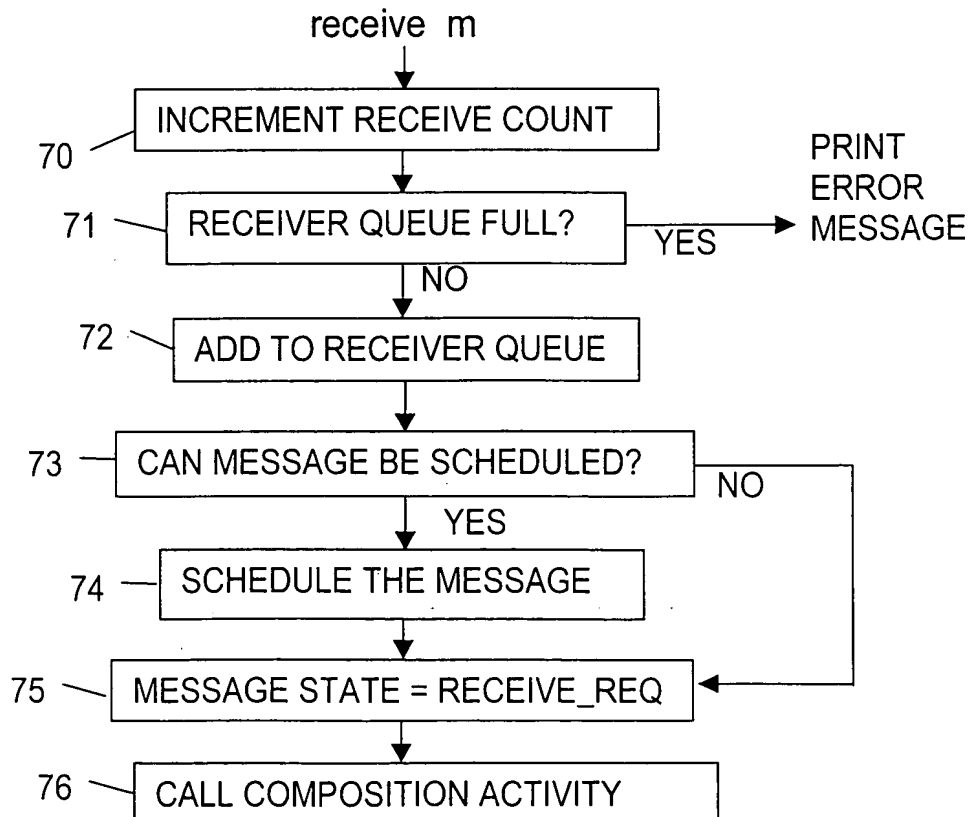


FIG.7



THIS PAGE BLANK (USPTO)

FIG.8

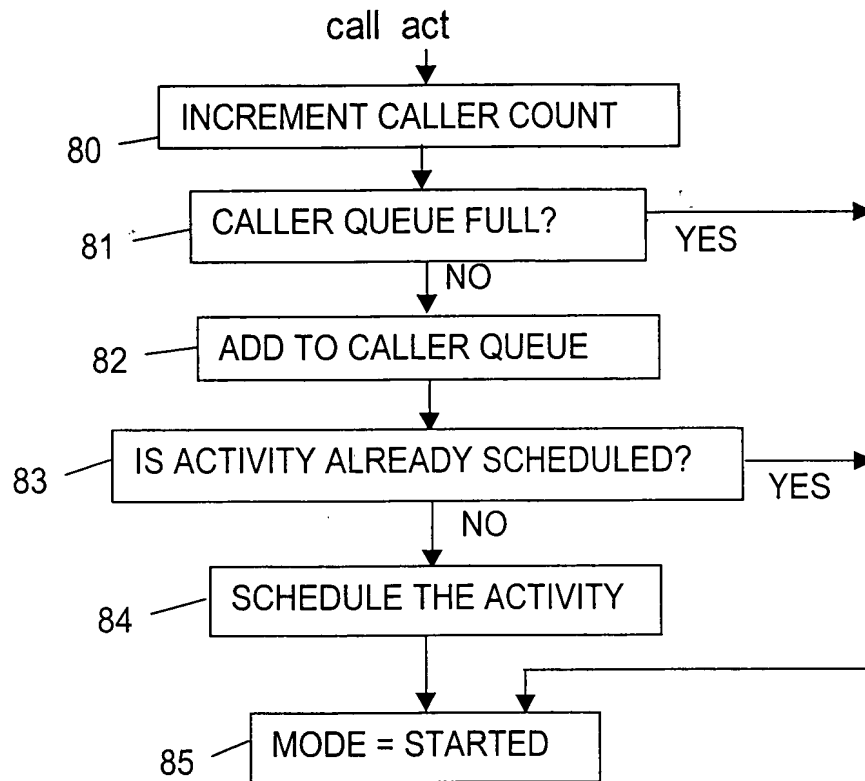
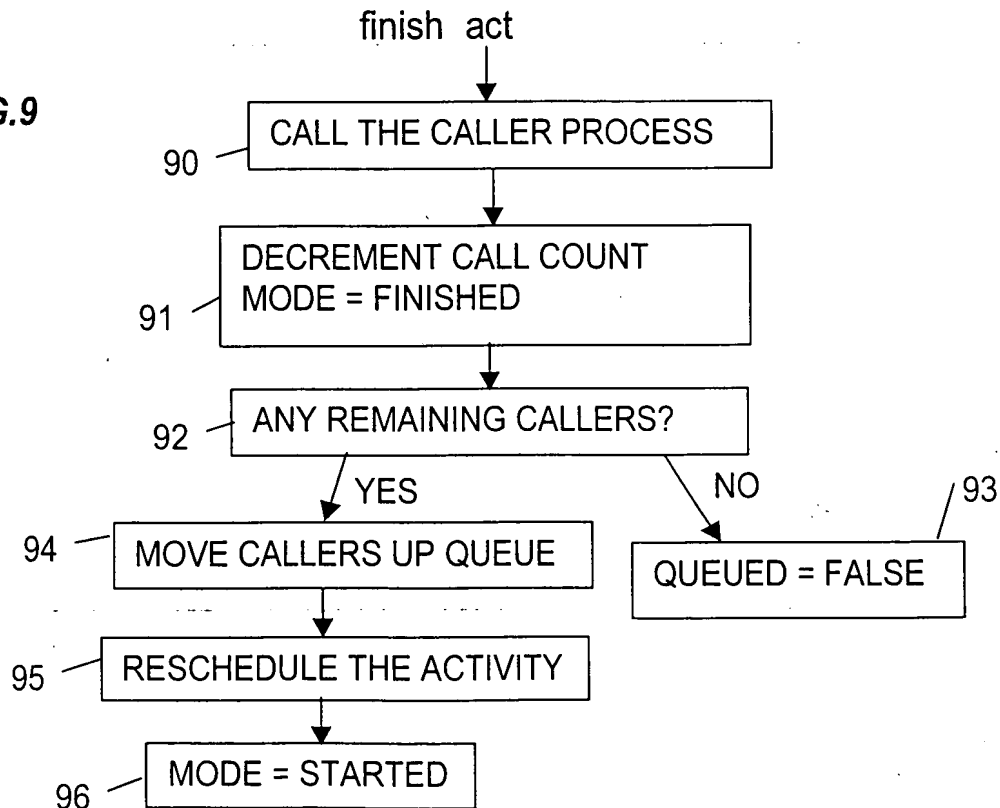


FIG.9



THIS PAGE BLANK (USPTO)

FIG.10

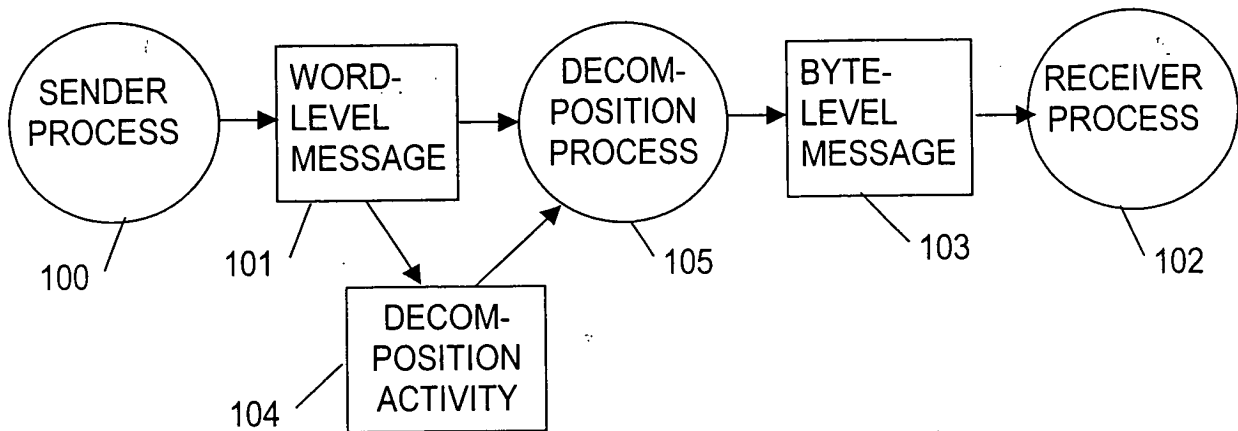
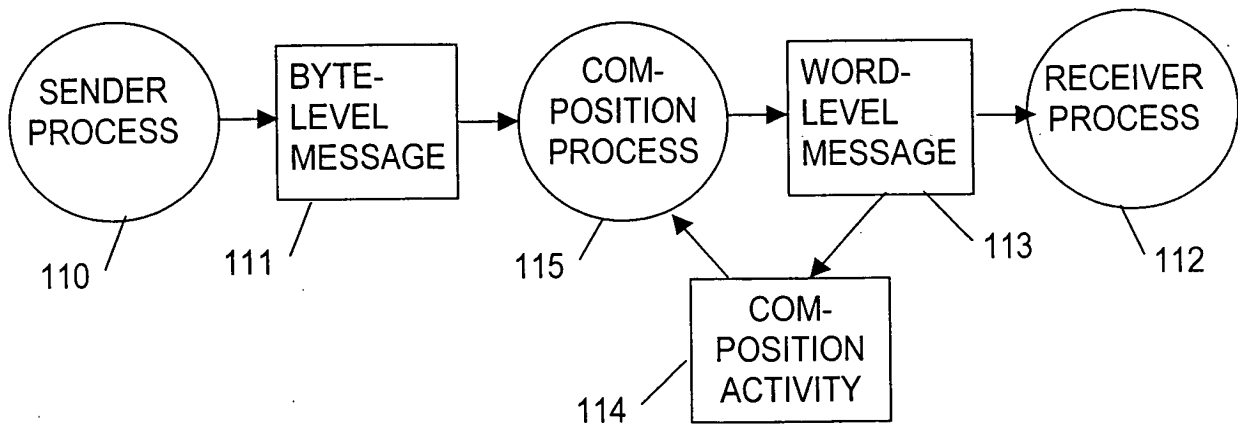


FIG.11



THIS PAGE BLANK (USPTO)

09761600

App # 09761600

6/14/04

Inventor Name Search-Palm, ACM, IEEE, Google

IEEE 6/16/04

+author:aho, +author:hopcroft

VHDL, VHDL+ interface

ACM 6/17/04

+simulation, +kernel, +queue, +event, +process, +signal, +digital system"

THIS PAGE BLANK (USPTO)